arXiv:1702.05198v1 [physics.plasm-ph] 17 Feb 2017

A Particle Down-Sampling Method with Application to Multifidelity Plasma Particle-in-Cell Simulations

D. Faghihi[1], V. Carey[2], C. Michoski[1], R. Hager[3], S. Janhunen[4], C. S. Chang[3], and R. D. Moser[1]

[1]Institute for Computational Engineering, University of Texas at Austin
[2]Department of Mathematics and Statistics, University of Colorado Denver
[3]Princeton Plasma Physics Laboratory, Princeton NJ
[4]University of Saskatchewan, Saskatoon, Canada

Abstract

In simulations of plasmas using particle-in-cell methods, it is often advantageous to resample the particle distribution function to increase simulation accuracy, to reduce the computing cost, or to avoid numerical instabilities. An algorithm for down-sampling the particles to a smaller number is proposed here. To minimize noise introduced by the down-sampling, the algorithm is designed to preserve important properties of the particle distribution, particularly low-order moments and local features of the distribution. Tests demonstrate that the proposed algorithm yields a more accurate representation of distributions than simple random down-sampling. The algorithmic approach can also be applied to other resampling problems.

Keywords: Particle-in-cell, particle resampling, distribution function moments, constrained optimization.

1 Introduction

Particle-in-cell (PIC) methods are widely used numerical methods for the solution of partial differential equations in a Lagrangian formulation. In particular, they are well-established for the solution of the coupled Boltzmann and Maxwell equations for the evolution of a plasma[1,2] and their gyrokinetic approximation[3]. In these methods, the numerical or marker particles can be viewed as mathematical samples from particle distribution functions describing the number density of atomic particle species making up the plasma (e.g. ions or electrons). The moments of the distribution function that appear as source terms in Maxwell's equations are then essentially computed using a Monte Carlo scheme. As is well known, Monte Carlo integration converges slowly (like

$1/\sqrt{N}$) with the number of samples, and is therefore a significant source of noise in PIC codes.

A wide variety of techniques have been employed to control and reduce these sampling errors, including importance sampling[4] and control variate techniques[5,6]. Resampling the particle distribution function is also a useful noise reduction technique[7]. The evolution of the particles can degrade the quality of the Monte Carlo sampling over time. This can occur because particles evolve away from the desired importance-sampled distribution. For example, in full-f gyrokinetic PIC tokamak simulations, low weight particles are loaded at the foot of the plasma pedestal and scrape-off layer to resolve important edge physics. However, during the simulation these low-weight particles mix with higher weight particles in the core, effectively spreading the marker particle weights and lowering the simulation efficiency in the critical pedestal and scrape-off layer regions. The weight spread can be ameliorated by periodic resampling to maintain the required marker particle weight spread in the edge regions. A related issue arises with the application of the collision operator in a PIC simulation. High energy atomic particles, often arising from non-Maxwellian processes, may be represented by few marker particles. The result is a noisy approximation of the collision operator at high energy. By sampling the high energy part of the distribution with many low-weight particles, and maintaining this state by periodically resampling the distribution, this problem can be alleviated.

Another application of resampling in PIC simulations arises when using control variate techniques as in δf gyrokinetic plasma simulators. In these methods, the particle weights evolve, which can lead to sample degeneration in which weights become concentrated on fewer and fewer particles as the simulation proceeds, resulting in a poor representation of the distribution. A similar phenomenon occurs in particle filtering[8–10]. Again, periodic resampling to rebalance the weights can address this problem.

Finally, resampling is needed when coupling between representations of differing fidelity. Examples of this include coupling between a gyrokinetic and full kinetic Boltzmann model, or between a spatially three-dimensional model and a spatially two-dimensional model. In this paper, we consider this latter case as an example. The mapping of a spatially three-dimensional PIC particle distribution to a spatially two-dimensional model allows the number of marker particles to be greatly reduced. This could arise, for example, when 3D kinetic PIC plasma simulations that resolve turbulent transport phenomena are reduced to 2D neoclassical transport evaluation using the particle distribution function distorted by turbulence.

In resampling a particle distribution in a plasma PIC simulation, it is important to preserve the low-order moments, because it is these low-order moments that are sources for the electromagnetic fields and transport. Further, the plasma evolution can lead to distributions that are far from equilibrium Maxwell-Boltzmann statistics, and these non-equilibrium features are also important to preserve in the resampling.

A general resampling algorithm may need to increase the number of marker

particles in a region of phase space and/or decrease the number of marker particles in a region of phase space. We call this up-sampling and down-sampling respectively. The algorithmic approach described in this paper is designed to allow the preservation of low-order moments and the structural features of the distribution for both up- and down-sampling. In this report, however, we consider the particular case of down-sampling as an example, because that is what is required for mapping a PIC kinetics solution from three spatial dimensions to two.

The remainder of the paper is outlined as follows: the multifidelity gyrokinetic application that motivates the current work is described in section 2. In section 3, the resampling algorithm is described, and example results are provided in section 4. Finally, concluding remarks are provided in section 5.

2 Multi-scale Gyrokinetic Simulation Using PIC Codes

The main motivation of the general purpose down-sampling algorithm presented here is to reduce noise in various applications of plasma simulation using PIC gyrokinetic codes. To provide context for this algorithmic development, the PIC algorithm and multifidelity approach are described briefly in this section.

The kinetic theory of collisional plasmas can be used to describe the behavior of gases and plasmas in terms of the motion of atomic particles (ions and electrons). Particle motion is represented by the evolution of the particle distribution functions $f(\mathbf{x}, \mathbf{v}, t)$, which measure the number density of atomic particles in a six-dimensional phase space $(\mathbf{x}, \mathbf{v})$. The evolution of the f's for various atomic species is described by a multi-species Boltzmann equation coupled to classical electrodynamics (i.e. Maxwell's equations). As an example, we consider electrostatic perturbations, by essentially replacing Maxwell's equations with Gauss's law.

For a given charged species this system can be written as:

$$\frac{\partial f}{\partial t} + \mathbf{v}\frac{\partial f}{\partial \mathbf{x}} + \frac{q}{m}(\mathbf{E} + \mathbf{v} \times \mathbf{B})\frac{\partial f}{\partial \mathbf{v}} = C(f) + S, \quad \nabla \cdot \mathbf{E} = -G, \tag{1}$$

where $\mathbf{E}$ is the electric field, $\mathbf{B}$ the magnetic field, C a collisional operator describing sub-grid Coulomb interactions, S a source and sink term, and G a forcing (net charge) for the electrostatics. Moreover, in strongly magnetized plasmas, such as fusion plasmas, these equations can be averaged over the gyrophase (viz. the cyclotron frequency) leading to a gyrocenter tracking equation for a charged ring over the relatively slow motion of particle gyrocenters, thus treating rapid particle orbits about magnetic field lines perturbatively. Transforming to the gyrocenter coordinate then results in a phase space reduction from six to five dimensions, and the resulting system of equations is referred to as the gyrokinetic equations.

While a significant simplification is achieved through the reduction of the full 6D equations to 5D gyrokinetic equations[3], simulating this 5D system over full-scale fusion reactor geometries is still a formidable task that requires carefully formulated

numerical approximation. One successful technique is to use PIC methods. In a PIC method, the distribution function f is sampled and the evolution of f is represented by the evolution of the samples in the 5D gyrokinetic phase space. The samples are (confusingly) called particles, marker particles, macro-particles, or sometimes super-particles. This is essentially a Lagrangian representation of the evolution of f, but the electromagnetic equations are solved on a Eulerian grid.

The objective is to represent f using many orders of magnitude fewer marker particles than physical atomic particles. For this reason, each marker particle carries a weight so that the sum of the weights over all the marker particles gives the number of physical particles. The marker particles can also be non-uniformly weighted for importance sampling, which allows more samples (particles) to be used to represent important parts of the phase space. For example, more particles with smaller weights can be used in low-density regions to provide a good representation of these regions. More particles can also be used in the high energy space when they become important.

The Lagrangian evolution equations for the marker particle (simply particles from now on) positions, velocities and weights depend on exactly how f is sampled. In the most straight-forward approach (called full-f), the distribution is simply sampled. However, this requires the largest number of particles to accurately represent f, and is therefore most computationally expensive. Other algorithms use a control-variate (δf) approach in which the difference between f and an ideal or simple distribution f_0 is sampled. There are a number of such δf approaches that differ in the details; see for example[11]. In the full-f representation, the particle weights do not evolve on the left-hand side operation of the Boltzmann equation (1) , but in the δf approaches, they do.

The details of a resampling or down-sampling algorithm for a full-f and δf sampling scheme may differ, because of the difference in representation. In the algorithm described here, we consider a full-f sampling representation of f.

One of the major challenges in the simulation of a tokamak plasma is the multi-scale nature of the problem, both in space and time. In a simulation that resolves both plasma turbulence and axisymmetric neoclassical dynamics, it will necessarily be extremely expensive to track slowly evolving phenomena. To ameliorate this issue, a spatially 3D particle distribution function distorted from Maxwellian may be down-sampled to a spatially 2D axisymmetric particle distribution to evaluate the neoclassical transport arising from the turbulence distortion. Because of the lower dimensionality and resulting reduced number of particles required to represent f, the cost of this lower-fidelity formulation is greatly reduced. This is an exemplary objective of the down-sampling algorithm proposed here. In particular, we have considered an example distribution function from the gyrokinetic PIC turbulence code XGC1[12,13].

3 Particle Down-sampling Strategy

In resampling the particles in a PIC code, it is important to preserve key features of the distribution they represent. First, several of the moments of the distribution f are physically meaningful quantities, and preserving them during resampling avoids introducing resampling error into them. For example, one might be concerned about preserving the following quantities:

$$\text{Number density} : n(\mathbf{x},t) = \int f(\mathbf{x},\mathbf{v},t)d\mathbf{v}, \tag{2}$$

$$\text{Momentum density} : \mathbf{m}(\mathbf{x},t) = m\int \mathbf{v}f(\mathbf{x},\mathbf{v},t)d\mathbf{v}, \tag{3}$$

$$\text{Kinetic energy density} : K(\mathbf{x},t) = \frac{1}{2}m\int |\mathbf{v}|^2 f(\mathbf{x},\mathbf{v},t)d\mathbf{v} \tag{4}$$

$$\text{Stress density} : T_{kl}(\mathbf{x},t) = m\int v_k v_l f(\mathbf{x},\mathbf{v},t)d\mathbf{v}. \tag{5}$$

In specific applications, other moment quantities might also be of interest. For example, in gyrokinetic tokamak plasma simulations, the canonical angular momentum is of particular interest.

The low-order moments with respect to position characterize the spatial distribution of the particles (and therefore mass and charge):

$$\bar{\mathbf{x}}(\mathbf{v},t) = \int \mathbf{x}f(\mathbf{x},\mathbf{v},t)d\mathbf{x}, \tag{6}$$

$$\mathcal{R}_{kl}(\mathbf{v},t) = \int x_k x_l f(\mathbf{x},\mathbf{v},t)d\mathbf{x}, \tag{7}$$

and should generally also be preserved. Preserving such moments avoids introducing resampling error into these quantities. Clearly then, it is of great value to preserve a set of low-order moments of the distribution during resampling, depending on the application.

Finally, the distribution may have non-Maxwellian features that are important to the physics of the system, and these should also be preserved. So, the objective is to down-sample the particles representing f while preserving moments and detailed features of the distribution.

When performing the down-sampling, we adjust the weight distribution of the particles. In some particle evolution models, the dynamics of the system require that the weights evolve, and in this case they may evolve away from an efficient representation of the distribution, as occurs in particle filtering algorithms[8–10]. Alternatively, the weights may have been set to support importance sampling in key regions of the distribution domain. As mentioned in the introduction, in a tokamak plasma simulation, one might want to populate the edge pedestal foot or scrape-off layer (SOL) of the plasma, with more particles because of the low plasma density and the importance

of edge physics to plasma containment. However, as particles evolve in phase space, they may evolve away from the plasma edge to the core plasma where fewer simulation particles with higher weight are needed. In this case, resampling can save computing time. Regardless of the reason, the resampling process provides an opportunity to reintroduce the desired sampling characteristics.

To generate a set of M^p macro-particles that preserve moments accurately, capture salient features of the distribution function, and allow for specified importance sampling, the proposed resampling techniques use a binning strategy. Both physical and configurational moments are preserved within each bin. The re-sampling strategy is outlined below.

In a spatially three-dimensional gyrokinetic PIC code, each particle carries with it six attributes, its weight w, spatial coordinates $\mathbf{x} = (r, z, \phi)$, and drift velocity $\mathbf{v} = (v_\parallel, v_\perp)$. In the down-sampling considered here, the distribution f is projected on a spatially two-dimensional space, assuming the solution is statistically homogeneous in the ϕ direction. The first step of the process is to divide the four-dimensional phase space $\boldsymbol{\zeta} = (r, z, v_\parallel, v_\perp)$ into bins and sort the particles into those bins. Depending on the application, it may be convenient to use the spatial discretization grid of the PIC method to define the bins in space, but in the examples considered here, both space and velocity are binned with a rectangular grid. There are then N^p particles sorted into N^b bins, with bin i containing N_i^p particles.

The objective of the down-sampling is to reduce the number of particles to $M^p < N^p$. To this end the target number of particles in each bin M_i^p needs to be determined. In general, a particle number density function $g(\boldsymbol{\zeta})$ with $\int_{\boldsymbol{\zeta}} g \, d\boldsymbol{\zeta} = M^p$ can be defined to allow importance sampling, so that then

$$M_i^p = \int_{\boldsymbol{\zeta}^i} g \, d\boldsymbol{\zeta}, \tag{8}$$

where $\boldsymbol{\zeta}^i$ is the $\boldsymbol{\zeta}$ domain of bin i. The down-sampling can then be accomplished by drawing M_i^p samples in each bin from the f distribution restricted to $\boldsymbol{\zeta}^i$. Assuming that $M_i^p \ll N_i^p$, this can be easily accomplished by drawing samples of the original particles with probability proportional to their weights using weighted sampling without replacement [14].

All that remains is to set the new weights $\tilde{w}_j^i$ of each down-sampled particle. In doing so, the moments discussed above will be preserved in each bin. The requirements

are then that for each bin i:

$$\sum_{j=1}^{M_i^p} \tilde{w}_j^i = \sum_{j=1}^{N_i^p} w_j^i, \tag{9}$$

$$\sum_{j=1}^{M_i^p} \tilde{\boldsymbol{\zeta}}_j^i \tilde{w}_j^i = \sum_{j=1}^{N_i^p} \boldsymbol{\zeta}_j^i w_j^i, \tag{10}$$

$$\sum_{j=1}^{M_i^p} (\tilde{x}_k \tilde{x}_m)_j^i \tilde{w}_j^i = \sum_{j=1}^{N_i^p} (x_k x_m)_j^i w_j^i, \tag{11}$$

$$\sum_{j=1}^{M_i^p} (\tilde{v}_k \tilde{v}_m)_j^i \tilde{w}_j^i = \sum_{j=1}^{N_i^p} (v_k v_m)_j^i w_j^i, \tag{12}$$

where k,$m \in \{1,2\}$, and the particles have been renumbered by bin (superscript i) to be sequentially numbered within the bin (subscript j) for both the original and down-sampled particles. In addition, because the down-sampling of the particles was weighted by weights w_j^i of the original particles, weights of the down-sampled particles $\tilde{w}_j^i$ should be uniform within each bin. However, this would not allow the constraints to be satisfied. So, instead, the weights are made as close to uniform over the bin as possible while satisfying the constraints (9-12). In particular, the new weights $\tilde{w}_j^i$ are set as

$$\{\tilde{w}_j^i\}_{(j \in 1,2,\ldots M_i^p)} = \underset{\tilde{w}_j^i \in \mathcal{C}^i}{\arg\min} \sum_{j=1}^{M_i^p} (\overline{\tilde{w}}^i - \tilde{w}_j^i)^2, \tag{13}$$

where

$$\overline{\tilde{w}}^i = \frac{1}{M_i^p} \sum_{j=1}^{N_i^p} w_j^i, \tag{14}$$

and $\mathcal{C}^i$ is the set of $\tilde{w}^i$ satisfying the constraints (9-12). Note that with a quadratic objective function and linear equality constraints in $\tilde{w}_j^i$, this is a straight-forward quadratic programming problem, which is solved using the algorithm of Goldfarb & Idnani [15], as implemented in the *quadprog* Fortran library [16]. Since the optimization problems for each bin is independent, they can easily be solved in parallel.

4 Numerical Examples

The feature-based resampling algorithm described in Section 3 is numerically implemented in a MPI-enabled Fortran code. While this method can be applied to resample any distribution, the primary aim of this implementation is to improve the accuracy of plasma particle simulations using gyrokinetic PIC codes. The example application

considered here is the down-sampling required to perform spatially two-dimensional neoclasical simulations based on characterization of the plasma distribution function obtained from a spatially three-dimensional simulation of the turbulence. In current implementations, this resampling is performed by randomly down-sampling particles over the entire domain of f, without preserving either first or second moments of position or velocity. The algorithm described here is expected to provide a significant improvement in the accuracy of the kinetics representation. To demonstrate the quality of the proposed down-sampling, we present the results of several test problems, including down-sampling of particles from an XGC1 turbulence solution.

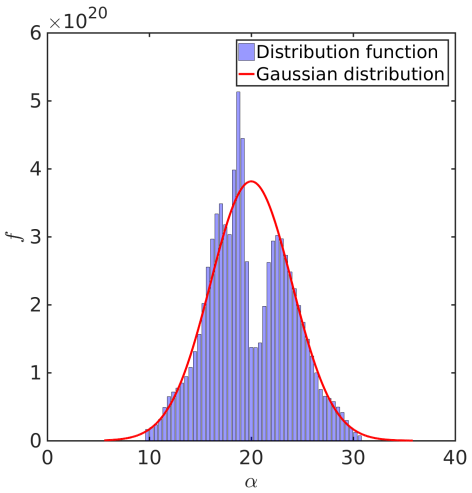

Figure 1: Initial particle distribution function with local features. In this case, total number of particles N^p is 50000. The solid red curve is a Gaussian.

To assess the ability of the resampling algorithm to preserve features of the distribution, consider the one-dimensional probability distribution shown in Figure 1. It might, for example, be considered to be the distorted marginal (reduced dimensional) distribution of one of the velocity components, in which case Maxwell-Boltzmann statistics would yield a Gaussian distribution, and the given distribution is clearly far from Gaussian. This distribution was represented with $N^p = 50000$ samples, and then down-sampled to $M^p = 1000$ and $M^p = 250$: down-sampling factors (DSF) of 50 and 200, respectively. Down-sampling was performed using a one-dimensional version of the algorithm proposed in Section 3 using $M^b = 50$ bins in which the sample density, and first and second moments are preserved. For comparison, global random down-sampling was also used. The results are shown in Figures 2 and 3. By construction, with the proposed algorithm, the histogram formed on the 50 down-sampling bins is identical to that of the original sample (compare Figures 2a and c with Figure 1), but the randomly down-sampled histograms show large variations from the original. Indeed, the sampling noise obscures the structure of the distribution; so much so that in

the DSF=200 case, the underlying structure of the distribution is not visible at all. By construction, the low- order moments of the distribution are preserved in the proposed approach, as shown in Figure 3, whereas relative errors in the moments with random down-sampling range from 2% to 10%.

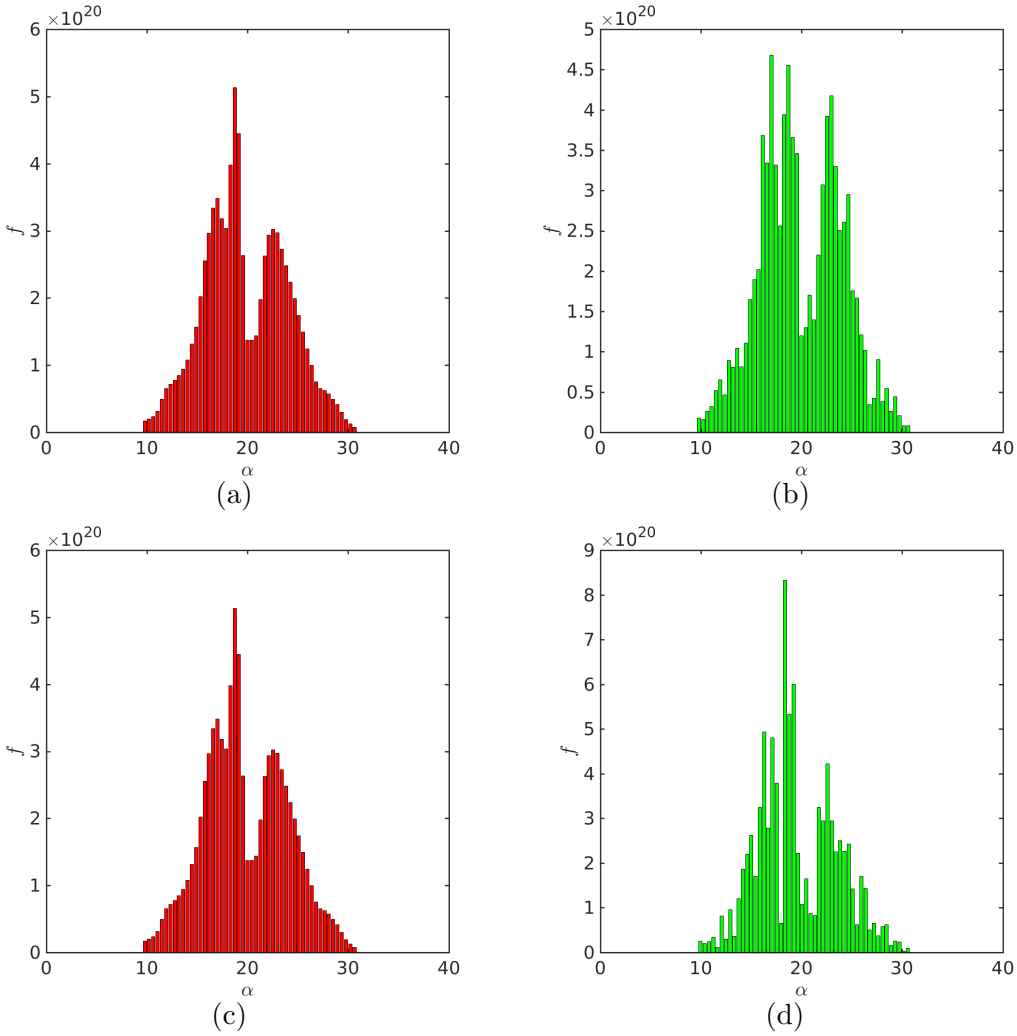

Figure 2: Local features preservation after particle down-sampling of the distribution in Figure 1 (number of bins = 50): (a) feature-based resampling method (DSF = 50); (b) random down-sampling (DSF = 50); (c) feature-based resampling method (DSF= 200); (d) random down-sampling (DSF = 200).

The proposed resampling algorithm was also tested on a particle distribution taken from a tokamak plasma simulation performed with the gyrokinetic PIC code XGC1. The data set consists of $N^p = 10^6$ particles, that were initially weighted to effect importance sampling, as discussed in Section 3. The particles have since

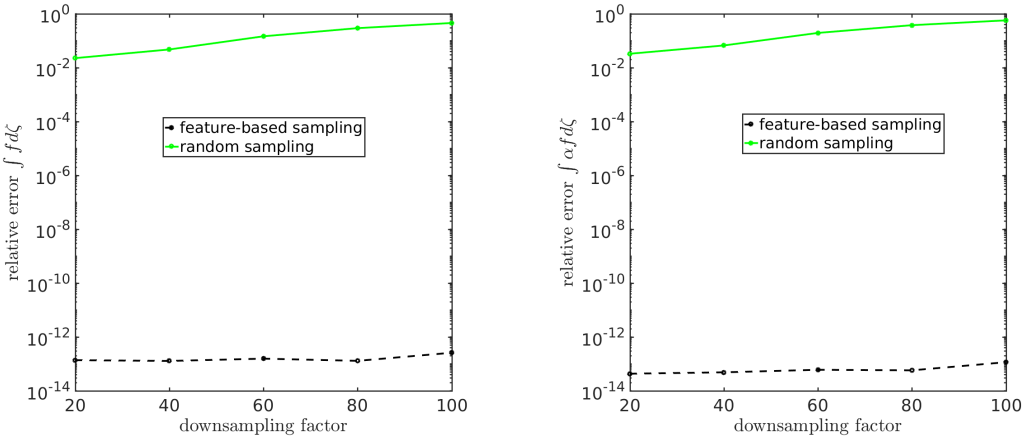

Figure 3: Relative error in global 0th and 1st moments due to feature-based/random down-sampling for different down-sampling factors (error in random down-sampling is the mean of 10^4 realizations).

been mixed through their evolution, so their weights no longer accomplish the desired importance sampling. For this test, to demonstrate the adjustment of weights as part of the resampling process, the importance weighting was eliminated by making $g \propto f$, where g is the marker particle distribution function in (8). In this case, the target number of particles M_i^p in each bin is simply proportional to the sum of the weights of the original particles in the bin. Specifically;

$$M_i^p = M^p \left(\frac{\sum_{j=1}^{N_i^p} w_j^i}{\sum_{j=1}^{N^p} w_j} \right). \tag{15}$$

In the gyrocentred XGC1 code, the particle velocity space is parameterized in terms of the normalized parallel velocity $\rho = v_\parallel / B$ and magnetic moment $\mu = v_\perp^2 / 2B$, where B is the magnitude of the magnetic field. For this case then, the four-dimensional down-sampled phase space is $\boldsymbol{\zeta} = (r, z, \rho, \mu)$. The $\boldsymbol{\zeta}$ domain was divided into bins, with 50 bins in each of the four phase-space directions. One complication that arises in sorting the particles into bins is that in some cases the target number of down-sampled particles M_i^p in a bin becomes too small to allow the moment constraints to be imposed. In this case, neighboring bins in the velocity space directions (ρ, μ) are merged to form larger bins with more particles, until M_i^p exceeds a specified minimum, which was set to 25 in the test conducted here.

The distribution of weights is shown in Figure 4 for the original particle and the down-sampled particles. Note that the original distribution is broad, with weights ranging over two orders of magnitude, while the resampled weights are much more narrowly distributed, varying by about 20% in magnitude. The resampled weights are

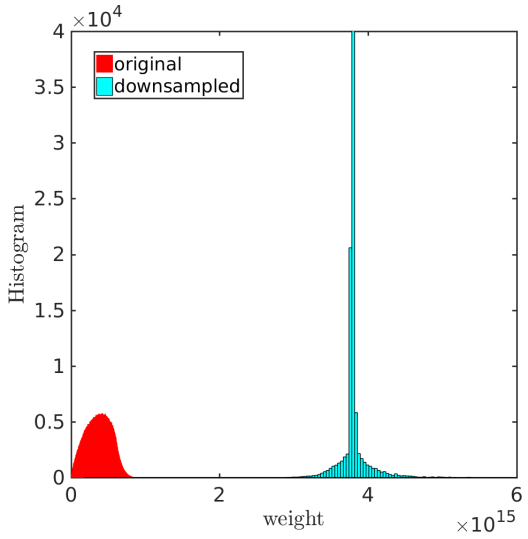

Figure 4: Histograms of particle weights from the original weights and the proposed feature-based resampling algorithm. In this case, $N^p = 10^6$, DSF = 20, and the number of bins in the ρ, μ, r and z directions is 50.

about 20 times larger in magnitude because there are 20 times fewer particles. Marginal distributions of velocity variables and of spatial locations are shown in Figure 5. As with the one-dimensional test, the original and down-sampled distributions have the same features.

The accuracy of the moment preservation in the feature-based resampling approach has been tested by performing the resampling while preserving just the zeroth moment, the zeroth and first moments and the zeroth, first and second moments as described in Section 3. This was done for a range of down-sampling factors, and the results are shown in Figures 6 and 7, along with those for a random down-sampling. The moment error for the random down-sampling are averaged over 10^4 sampling realizations. In all cases shown in Figures 6 and 7 the relative error of random down-sampling increases with down-sampling factor (decreasing number of down-sampled particles). In all cases, when the moments are constrained, the relative errors in the down-sampled moments are of order 10^{-10}. When only some moments are constrained, the errors in those that are not are significantly lower than in the random case. Presumably this is because of the binned down-sampling algorithm, which by itself should reduce the moment errors relative to purely random down-sampling.

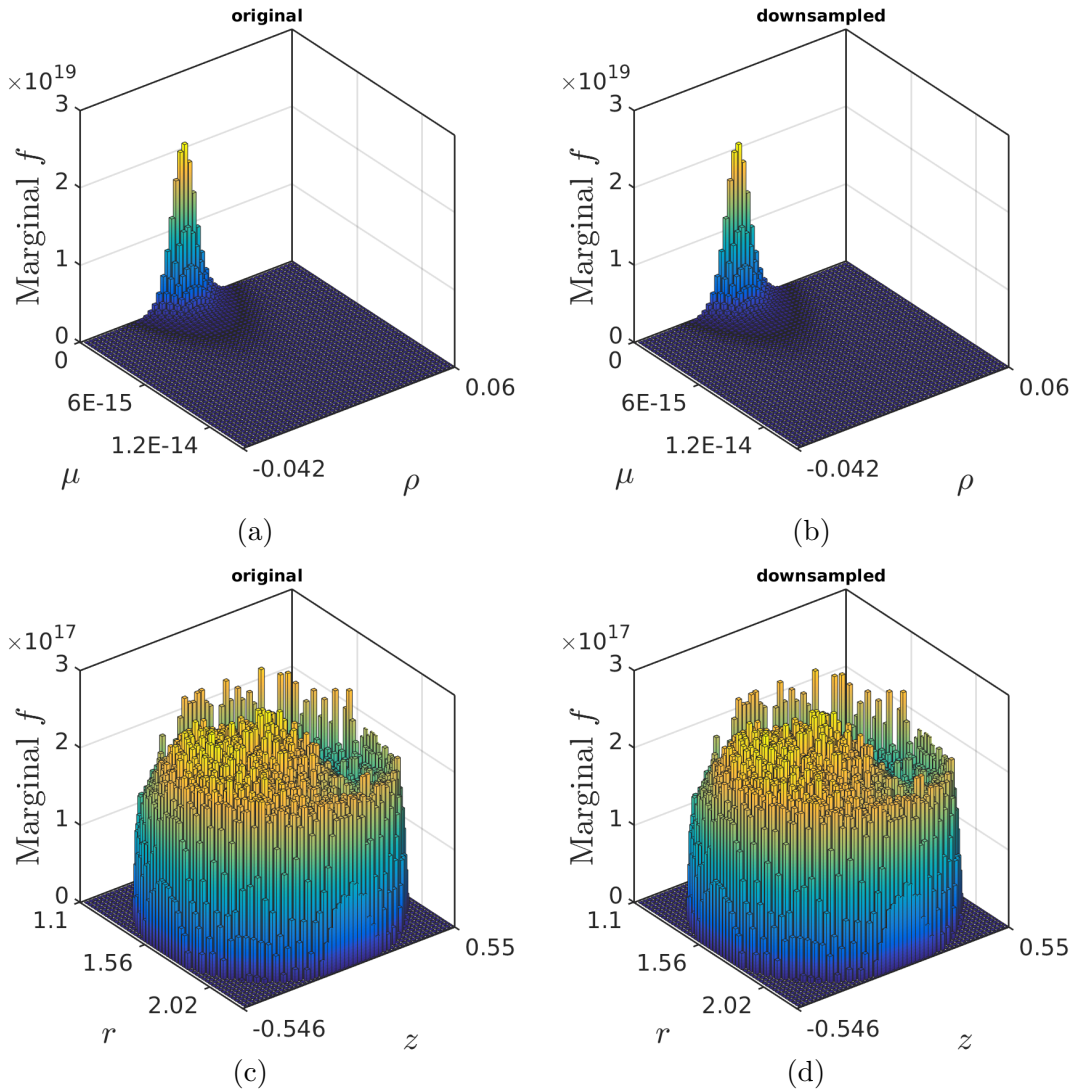

Figure 5: Marginal 2D bar plots of distribution functions before and after particle down-sampling (a and b) along ρ and μ; (c and d) along r and z. In this case, $N^p = 10^6$, DSF = 20, and the number of bins in the ρ, μ, r and z directions is 50.

5 Summary and Conclusions

The resampling algorithm presented here was designed to reduce the noise introduced by resampling particle distribution functions in a PIC code while preserving important features of the distribution. By performing the down-sampling in bins and constraining the low-order moments computed in each bin, the algorithm preserves features of the particle distribution that are important to the physical dynamics of the system. The

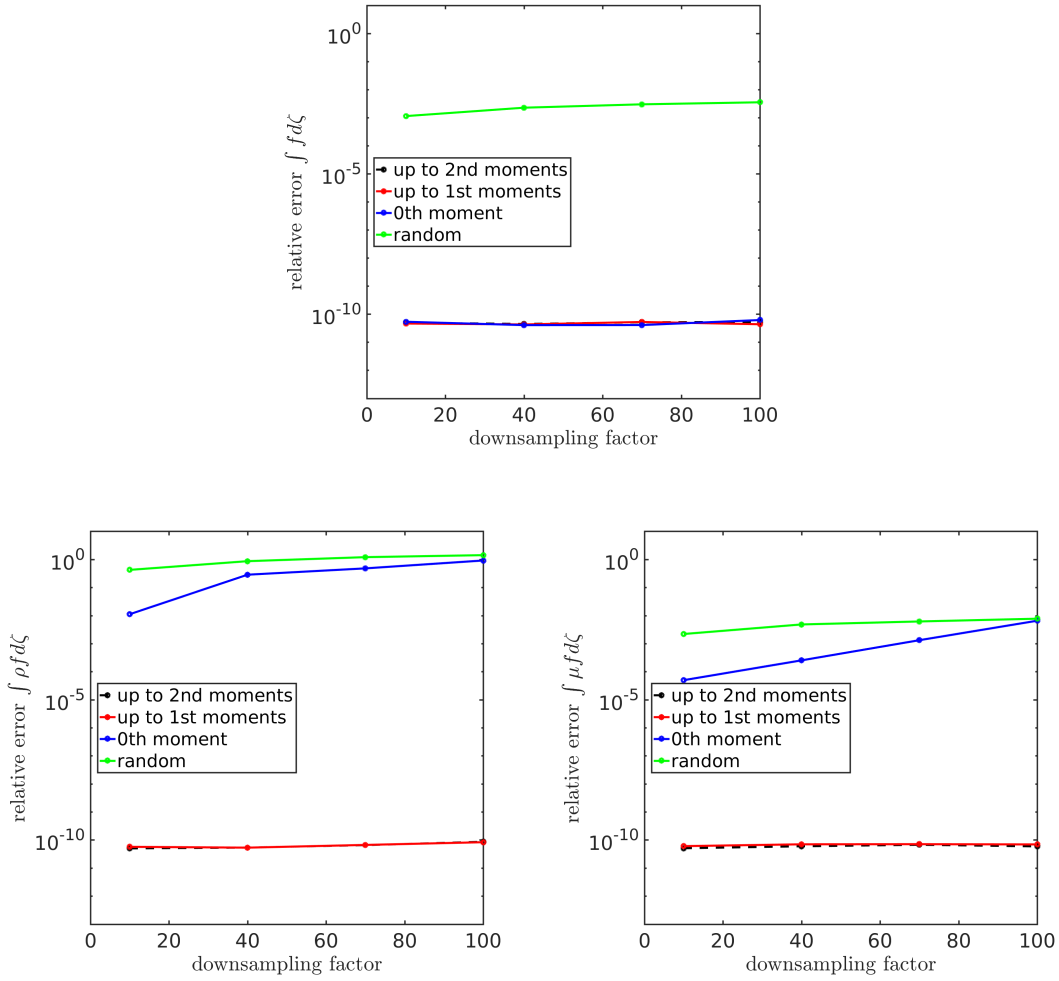

Figure 6: Relative error in global 0th and first moments due to feature-based/random down-sampling for different down-sampling factors. In this case, $N^p = 10^6$ and the number of bins in the ρ, μ, r and z directions is 10.

binning also enables the introduction or adjustment of importance sampling. While the algorithm was developed in support of PIC simulations of tokamak plasma, it is widely applicable to samples of general distributions, provided the dimension is not so large as to make binning impractical.

The general strategy of resampling a distribution using bins can also be applied in other contexts. For example, in the gyrokinetic PIC application discussed here, the up-sampling required to transfer a solution from a low-fidelity model to a higher fidelity model is also of interest, and can be pursued in a similar way, with similar benefits. In

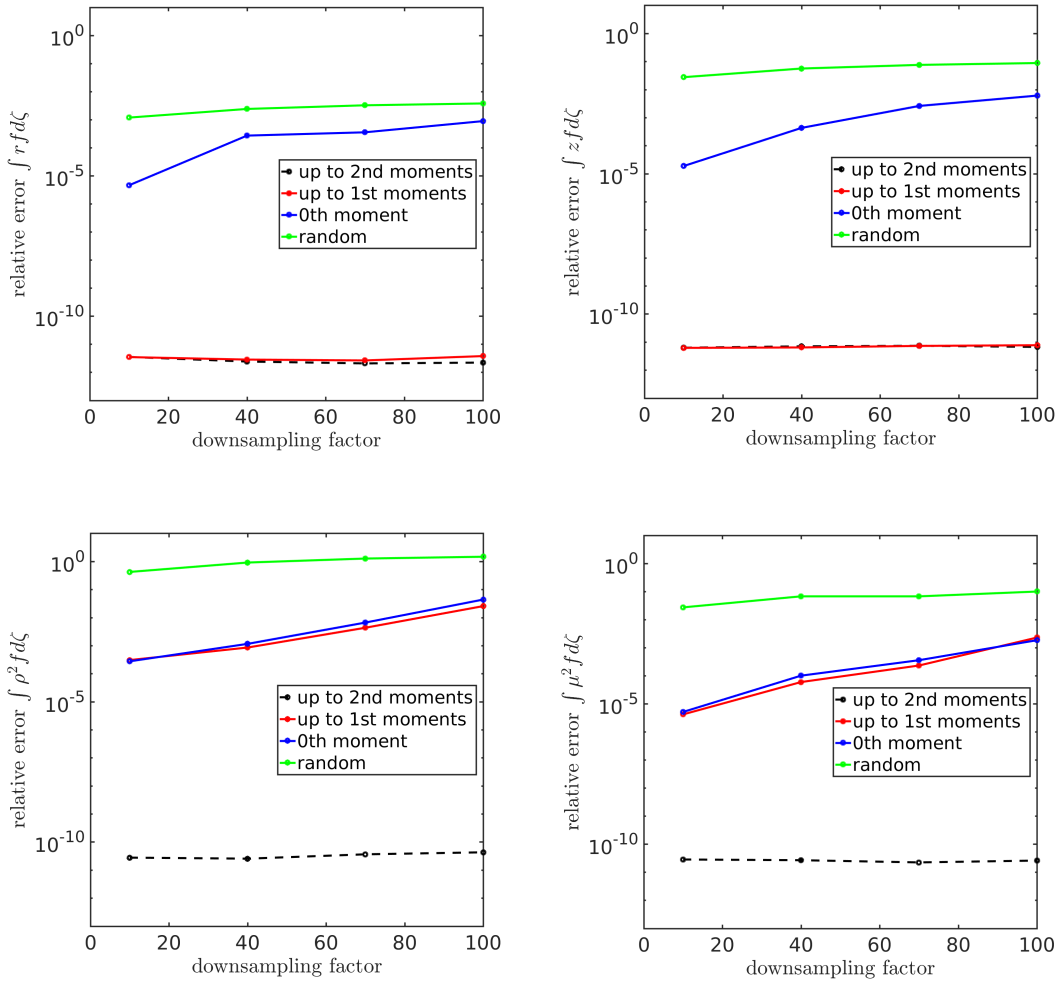

Figure 7: Relative error in global first and second moments due to feature-based/random down-sampling for different down-sampling factors. In this case, $N^p = 10^6$ and the number of bins in the ρ, μ, r and z directions is 10.

each bin, samples would need to be drawn from a distribution that matches the low-order moments of the original particles. The same algorithm would also generally be useful for resampling to adjust importance sampling weights, or to re-balance weights to avoid degeneration as occurs in particle filters.

Finally, there are two generalizations of the algorithm described here that could be particularly useful in PIC codes. First, the binning of the spatial location of the particles could be based on the spatial grid used in the Eulerian part of the calculation. This would be computationally convenient, since particles must usually be sorted into

the grid cells as part of the transfer of information to the Eulerian calculation. Second, the operator that maps the particle configuration in each cell onto the Eulerian grid for use in the Eulerian calculation could be used as a constraint in the resampling, in addition to other low-order moments. In this way, the resampling would not introduce any change in the Eulerian solution.

Acknowledgments

The work reported here was supported by the Department of Energy under grants DE-SC0008454 and DE-AC02-09CH11466. The authors acknowledge the Texas Advanced Computing Center (TACC, http://www.tacc.utexas.edu) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this paper.

References

[1] Charles K Birdsall and A Bruce Langdon. *Plasma physics via computer simulation*. CRC Press, 2004.

[2] Roger W Hockney and James W Eastwood. *Computer simulation using particles*. CRC Press, 1988.

[3] TS Hahm. Nonlinear gyrokinetic equations for tokamak microturbulence. *Physics of Fluids (1958-1988)*, 31(9):2670–2673, 1988.

[4] Ahmet Y Aydemir. A unified Monte Carlo interpretation of particle simulations and applications to non-neutral plasmas. *Physics of Plasmas (1994-present)*, 1(4):822–831, 1994.

[5] Eric Sonnendrücker, Abigail Wacher, Roman Hatzky, and Ralf Kleiber. A split control variate scheme for pic simulations with collisions. *Journal of Computational Physics*, 295:402–419, 2015.

[6] R. Kleiber, R. Hatzky, A. Knies, K. Kauffmann, and P. Helander. An improved control-variate scheme for particle-in-cell simulations with collisions. *Computer Physics Communications*, 182(4):1005 – 1012, 2011.

[7] Roman Hatzky, Trach Minh Tran, Axel Knies, Ralf Kleiber, and Simon J. Allfrey. Energy conservation in a nonlinear gyrokinetic particle-in-cell code for ion-temperature-gradient-driven modes in -pinch geometry. *Physics of Plasmas*, 9(3):898–912, 2002.

[8] Neil J Gordon, David J Salmond, and Adrian FM Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. In *IEE Proceedings F-Radar and Signal Processing*, volume 140, pages 107–113. IET, 1993.

[9] Genshiro Kitagawa. A Monte Carlo filtering and smoothing method for non-gaussian nonlinear state space models. In *Proceedings of the 2nd US-Japan joint seminar on statistical time series analysis*, pages 110–131, 1993.

[10] Genshiro Kitagawa. Monte Carlo filter and smoother for non-gaussian nonlinear state space models. *Journal of computational and graphical statistics*, 5(1):1–25, 1996.

[11] S. Ku, R. Hager, C.S. Chang, J.M. Kwon, and S.E. Parker. A new hybrid-lagrangian numerical scheme for gyrokinetic simulation of tokamak edge plasma. *Journal of Computational Physics*, 315:467 – 475, 2016.

[12] C.S. Chang and S. Ku. Spontaneous rotation sources in a quiescent tokamak edge plasma. *Physics of Plasmas*, 15(6), 2008.

[13] C. S. Chang, S. Ku, P. H. Diamond, Z. Lin, S. Parker, T. S. Hahm, and N. Samatova. Compressed ion temperature gradient turbulence in diverted tokamak edge. *Physics of Plasmas*, 16(5):056108, 2009.

[14] Chak-Kuen Wong and Malcolm C. Easton. An efficient method for weighted sampling without replacement. *SIAM Journal on Computing*, 9(1):111–113, 1980.

[15] Donald Goldfarb and Ashok Idnani. A numerically stable dual method for solving strictly convex quadratic programs. *Mathematical programming*, 27(1):1–33, 1983.

[16] A.T. Berwin. quadprog, a FORTRAN library for solving quadratic programming problems. `https://github.com/badgerodon/quadprog`, 2013.